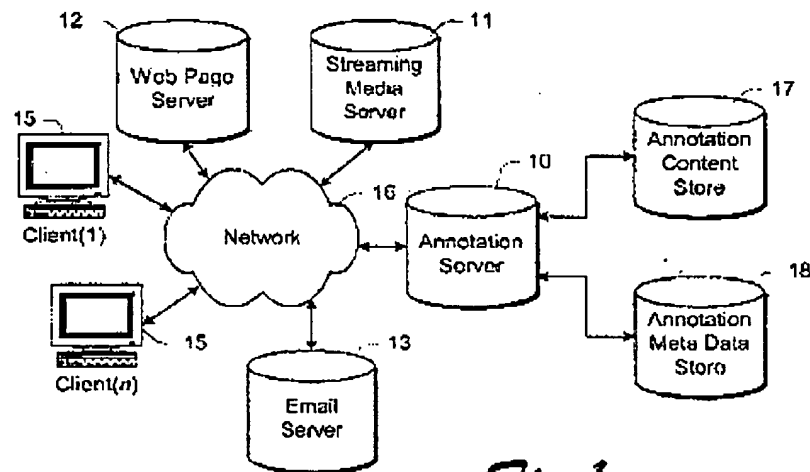


ATTY. DKT NO. 15

PATENT

controls the delivery of the annotations, which are stored in an annotation content store 17. Much of the patent application is devoted to explaining how end users at client machines 15 create and upload annotations to the annotation server 10. That aspect of the Gupta et al. disclosure does not appear relevant to the subject invention, however.

*Fig. 1*

As can be seen above, a second data store 18 stores "meta data" associated with a given annotation. The meta data typically is in the form shown in Figure 5 of the application and includes, e.g., information identifying the author of the annotation, timestamps identifying to which portions of the media stream the annotation applies, the creation time, information identifying the annotation or related annotations, information identifying which versions of the media stream the annotation applies, and so forth. As illustrated in Figure 11, once a given media stream is being received at a client machine 15, the end user can obtain a previously-created annotation stored in the annotation content store 17. To this end, the annotation server identifies the media characteristics of the media stream, fetches the annotation, and serves that annotation to the requesting end user client machine. During this process, information (e.g., timestamp data) in the annotation meta data store presumably is accessed and used by the annotation server to ensure that the annotation is synchronized with the media stream to which it applies.

ATTY. DKT NO. 15

PATENT

Because the "annotation server" is the only device in Gupta et al. that applies information from the meta data store 18, it appears then that the Examiner is equating the "annotation server" 10 with the described "content servers" of the "content delivery network." (See, e.g., claim 17). Moreover, the Examiner also appears to be equating the "annotation" itself with the "given content to be delivered over the content delivery network" and, further, that the information obtained from the meta data store 18 is the described "content control requirement." These contentions are misplaced.

Respectfully, the annotation system described in Gupta et al. is not a "content delivery network" comprising "a set of content servers," which content server (in the plural) are used "on behalf of participating content providers" who have identified "given content to be delivered over the content delivery network." (See, e.g., claim 17). Indeed, as is evident from Figure 1 reproduced above, in Gupta et al. the content itself (typically a media stream) is obtained from the streaming media server 11 (and possibly the web page server 12), but there is one, and only one, "annotation server 10" for serving the annotations. Thus, even if one were to fairly construe an annotation author as a "content provider" and an "annotation" as "given content," which Applicants do not concede, the claims require the invention to be implemented within the context of a "content delivery network" having "a set of content servers." There is no disclosure or suggestion in Gupta et al. to use a distributed set of annotation servers. In this regard, the Examiner is directed to the limitation in each of independent claims 17 and 33 requiring that the "metadata for the given piece of content" be communicated "to the set of content servers." Gupta et al. have no need for this step because they have just one annotation server, not a distributed set of such servers. Because Gupta et al. do not have multiple annotation servers and no need to communicate annotation meta data anywhere, they also fail to meet the limitations in dependent claims 19 (communication by header), 20 (communication by configuration file), 21 (communication by configuration file provisioned via an extranet application), 26 (communication), 27 (communication by a configuration file), 32 or 34 (communication by one of: a request string, a header and a configuration file), or 35 (communication by one of: a request string, a header and a configuration file, as provisioned via an extranet application).

ATTY. DKT NO. 15

PATENT

Each of independent claims 17, 25 and 33 further require that the “given content control requirement specified in the metadata” be applied “at the given content server” prior to serving the given piece of content. Once again then, even if one were to fairly construe the annotation as the “given piece of content,” the actual limitation requires that the “given content control requirement specified in the metadata” be applied at a content server that has been selected or identified from the “set of content servers” in the content delivery network. This requirement cannot be met in Gupta et al., as there is only one annotation server and, thus, no selection or identification of a “given” one selected or identified from a set of such servers.

It is also doubtful that one of ordinary skill in the art would read “given content control requirement” to reach the information that is stored in the annotation meta data store 18. In particular, the written description describes the CDN server metadata as a “set of control options and parameters that determine how a CDN content server will handle a request for an object” (see, e.g., text at page 8, line 21; page 2, lines 26-31 through page 3, line 2, emphasis supplied). Examples of such content control requirement metadata are provided, for example, on page 8, lines 25-30 and throughout the written description. For the reasons set forth above, the lone annotation server 10 is not a content delivery network (CDN) content server, i.e., one of a set of content servers to which participating content providers offload their content for delivery to requesting end users.

Turning to the dependent claim rejections, the Examiner errs in his conclusion (Office action, paragraph 9) that Gupta et al. teaches communication metadata in a configuration file. Paragraph [82] cited by the Examiner merely describes how an end user computer establishes a connection to the annotation server. There is nothing in this paragraph that describes communicating any annotation meta data. Likewise, paragraph [91] merely describes that the annotation – not the meta data – can be sent to some intended recipient.

As per claim 21, it is respectfully submitted that “email” is not an “extranet application” but, nevertheless, it should be appreciated that an email is not a “configuration file” and there is nothing in Gupta et al. that suggests using an “extranet application,” such

ATTY. DKT NO. 15

PATENT

as a secure content provider customer portal, to generate a configuration file, which is then communicated to "a set of content servers."

Respectfully, regarding claim 22, the Examiner is also incorrect in equating the "given authentication method" limitation with the subject matter of paragraph [60] or the "given access control method" with the subject matter of paragraph [91]. Paragraph [60] simply states that a user who creates an annotation is identifiable as an author based on client logon information; paragraph [90], as described above, explains how an end user can send an annotation, not annotation meta data, to a target recipient. The claim, in contrast, is referring to applying an authentication-based or access control-based "content control requirement." Thus, even if the annotation server is construed as the "content server" and the "annotation" as the "given content," there is nothing in Gupta et al. (especially the annotation handling routine of Figure 11) that indicates some authentication or access control method applied at the annotation server itself.

Regarding claims 23-24, the cited references (paragraphs [47] and [91]) do not distinguish in any way the meta data as "request" or "response" meta data.

The Examiner, however, is correct in concluding that Gupta et al. do not explicitly teach the steps of "having a participating content provider associate a content provider domain or subdomain with a domain managed by a content delivery network service provider" or "resolving a client query to the content provider domain or subdomain to an IP address of a given content server in the set of content servers using the domain managed by the content delivery network service provider." (See Office action, at page 3). Nevertheless, these limitations are said to be provided by the secondary reference, Shobatake et al. This contention is incorrect, for the following reasons.

In the first instance, Shobatake et al. is non-analogous art. (See, MPEP §2141.01(a)). The patent has nothing to do with Internet content delivery generally, content delivery networks in particular, or methods of controlling how content is served from content servers. In contrast, the patent describes how mobile devices (e.g., cellphones, PDAs, etc.) can be used across multiple, disparate communications platforms. In this system, mobile device users desire "the ability to communicate with anyone, irrespective of the network in which the called party (or callee) is located). Because each of

ATTY. DKT NO. 15

PATENT

the mobile device networks have different protocols and technologies, no direct connection is possible.” (Column 2, lines 42+). Shobatake et al. address this problem by providing a unified mobility manager mobility that enables a mobile device to register with its “home database” even as the device moves within and across foreign networks.

One of ordinary skill in the art of “content delivery” would not look to the mobile device telecommunications art to address the problems addressed by the present invention. As described in page 3 of the written description, the present invention deals generally with the problem of how content providers who are using a “content delivery network” can set up and ensure that given content control requirements are applied to the content being served by the CDN service provider’s content servers. This is a completely different problem from that addressed by Shobatake et al., who are addressing the problem of how to recognize a mobile device user as that user moves across disparate telecommunications networks. This significant difference is further evidenced by the classification of Shobatake et al. in class 455, which is unrelated to Internet or Web-based enabling technologies and infrastructure.

More specifically, Shobatake et al. also do not disclose Internet domain name service (DNS) functionality, let alone the specific requirements of having a content provider “associate a content provider domain or subdomain with a domain managed by a content delivery network service provider” or using that domain to resolve something else, namely, a “client query to the content provider domain or subdomain.” In other words, the claims require the resolution of the “domain managed by [the] content delivery network service provider” where the client query is directed to some other domain or subdomain. This concept is nowhere taught or suggested by Shobatake et al, who merely disclose a set of alias databases 305-308; these databases, however, merely store “identification information relating to mobile terminals on other platforms.” (Column 5, lines 34-36). Thus, e.g., alias database 305 might store a cellular telephone user’s alias so that the mobility manager can recognize the mobile device user even as he or she moves into some foreign network. The claim limitations are not related to storing alias information per se; rather, as noted above, the claims require a specific type of alias (a “domain managed by a content delivery network service provider”) the processing of that alias (in lieu of some

ATTY. DKT NO. 15

PATENT

other Internet domain or subdomain actually associated with the client query), as well as the further requirement that the domain processing identify a given content server from a "set of content servers" in a content delivery network.

Because Shobatake et al. is non-analogous art, it cannot be combined with Gupta et al. absent some express teaching, suggestion or motivation to select the teachings of the separate references and combine them to produce the claimed combination. With all due respect, the Examiner's contention in this regard (obvious to combine "because user terminals querying a content provider domain with a domain managed by a content delivery network service provider ... would enhance the capabilities of Gupta by allowing for handling of communications in a unified manner") does not appear to make sense. In the first instance, as noted above, the claims do not say that "user terminals" query a "content provider domain with a domain managed" by the CDN service provider, as the Examiner states. Rather, the claims require that a given client query is to "the content provider domain or subdomain" and that such query be handled using not the "content provider domain or subdomain" in the query but, rather, the "domain managed" by the CDN service provider. Moreover, the "user terminals" in Shobatake et al. are mobile devices, such as cellphones or PDAs, and the problem of mobility management does not have anything to do with how Gupta et al.'s annotation system might be modified to derive the subject matter (as a whole) of any pending claim.

For these reasons, the obviousness rejection should be withdrawn.

The Original Oath

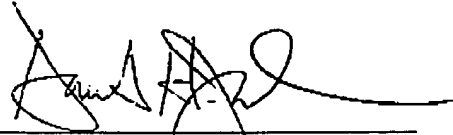
In view of the Examiner's comment in paragraph 3, the undersigned is in the process of obtaining a Substitute Declaration from each named inventor. To that end, a Substitute Declaration is submitted herewith for each of the inventors John Josef Kloninger and Philip A. Lisiecki. The Substitute Declarations for the other inventors will be filed shortly.

ATTY. DKT NO. 15

PATENT

Respectfully submitted,

By:

  
\_\_\_\_\_  
David H. Judson, Reg. No. 30,467

15950 Dallas Parkway  
Suite 225  
Dallas, Texas 75248  
(972) 385-2018